arXiv:1402.5568v1 [stat.ME] 23 Feb 2014

REGULARIZED BLOCK TOEPLITZ COVARIANCE MATRIX ESTIMATION VIA KRONECKER PRODUCT EXPANSIONS

Kristjan Greenewald, Alfred O. Hero III

University of Michigan, Ann Arbor

ABSTRACT

In this work we consider the estimation of spatio-temporal covariance matrices in the low sample non-Gaussian regime. We impose covariance structure in the form of a sum of Kronecker products decomposition [1, 2] with diagonal correction [2], which we refer to as DC-KronPCA, in the estimation of multiframe covariance matrices. This paper extends the approaches of [1] in two directions. First, we modify the diagonally corrected method of [2] to include a block Toeplitz constraint imposing temporal stationarity structure. Second, we improve the conditioning of the estimate in the very low sample regime by using Ledoit-Wolf type shrinkage regularization similar to [3]. For improved robustness to heavy tailed distributions, we modify the KronPCA to incorporate robust shrinkage estimation [4]. Results of numerical simulations establish benefits in terms of estimation MSE when compared to previous methods. Finally, we apply our methods to a real-world network spatio-temporal anomaly detection problem and achieve superior results.

1. INTRODUCTION

Multivariate processes are ubiquitous in signal processing. Our goal is to learn correlations between (possibly non Gaussian) variables both within a time instant (frame) and across time. The covariance for multivariate temporal processes manifests itself as multiframe covariance, i.e., the covariance not only between variables in a single frame, but also between variables in a set of nearby frames. If each frame contains p variables, then the covariance is described by a pT by pT matrix:

$$\mathbf{\Sigma}_t = \mathrm{Cov}\left[\{\mathbf{X_n}\}_{n=t-T}^{t-1}\right], \tag{1}$$

where $\mathbf{X}_n$ denotes the p variables in the nth frame.

We focus on large p small n covariance estimation, that is, the high dimensional regime where the number of variables exceeds the number of training samples available to learn the covariance. For $p \geq n$, the standard sample covariance matrix (SCM) is the maximum likelihood covariance estimator. When n is on the order of p or smaller, however, it is well known that the SCM has a very undesirable poorly conditioned eigenstructure which results in the poor conditioning of the inverse of the SCM, leading to poor estimates of the inverse covariance ($\mathbf{\Sigma}^{-1}$). Accurate estimation of the inverse covariance is critical for many applications, including graphical modeling, prediction, anomaly detection, and classification.

One approach for addressing this problem is shrinkage estimation, which uses a weighted average of the sample covariance and a deterministic covariance, often chosen as a diagonal matrix called the shrinkage target. This can significantly improve the accuracy of the inverse of the estimate due to the improved eigenspectrum. The optimal amount of shrinkage, i.e., the shrinkage weights, is often estimated as the minimizer of the expected L2 estimation error [5, 3]. When the samples are heavy-tailed non-Gaussian an improved shrinkage estimator can be constructed by sphering the samples and using Tyler ML iterations, making the covariance estimate more robust to outliers [4].

In this work, we develop shrinkage estimators of covariance matrices having spatiotemporal structure. Estimation of spatio-temporal covariance matrices based on reducing the number of parameters via a truncated sum of Kronecker products representation, which we call KronPCA, is discussed in [1, 2], and asymptotic performance analysis [1] predicts significant gains in estimator MSE in the high dimensional regime as $n, p \to \infty$. However, the empirical performance suffers from poor conditioning of the estimate.

To compensate for this poor conditioning of KronPCA we introduced diagonally corrected Kronecker PCA (DC-KronPCA) [2]. This paper extends [2] in two directions. First, the MSE analyses of the shrinkage estimator in [3, 4] is used to estimate the shrinkage parameter of DC-KronPCA in the Gaussian and heavy-tailed regimes, respectively. Second, when the measurements are known to be temporally stationary, DC-KronPCA is formulated with a temporal block Toeplitz constraint [6].

To illustrate the proposed DC-KronPCA based methods, we consider a real world application: spatiotemporal anomaly detection in sensor networks. We incorporate temporal information into the spatio-temporal covariance as opposed to the single frame covariance based method of [4].

The rest of this paper is organized as follows: in Section 2, we review DC-KronPCA and present our block Toeplitz modification of it. Our methods for standard and robust shrinkage of the KronPCA estimate are given in Section 3, and numeri-

cal simulation results are presented in Section 4. We discuss an application to sensor network anomaly detection in Section 5, and conclude in Section 6.

2. BLOCK TOEPLITZ KRONECKER PCA

In this section, we consider the regularization of the sample covariance via decomposition into a sums of (space vs. time) Kronecker products representation (KronPCA) [1, 2]. We thus develop a method for learning block Toeplitz KronPCA covariances.

Let $\mathbf{X}$ be a $p \times T$ matrix with entries $\tilde{x}(m,t)$ denoting samples of a space-time random process defined over a p-grid of space samples $m \in \{1,\ldots,p\}$ and a T-grid of time samples $t \in \{1,\ldots,T\}$. Let $x = \mathrm{vec}(\mathbf{X})$ denote the pT column vector obtained by lexicographical reordering. Define the $pT \times pT$ spatiotemporal covariance matrix $\mathbf{\Sigma} = \mathrm{Cov}[x]$. Following [2, 1, 7], we define the DC-KronPCA model

$$\mathbf{\Sigma} = \left(\sum\nolimits_{i=1}^{r} \mathbf{T}_i \otimes \mathbf{S}_i\right) + \mathbf{I} \otimes \mathbf{U}, \tag{2}$$

where $\mathbf{T}_i$ are $T \times T$ Toeplitz matrices (the temporal Kronecker factors), $\mathbf{S}_i$ are $p \times p$ matrices (the spatial Kronecker factors), $\mathbf{U}$ is a $p \times p$ diagonal matrix.

Following the approach of [1], we propose to fit the model (2) to the sample covariance matrix $\mathbf{\Sigma}_{SCM} = n^{-1}\sum_{i=1}^{n}(x_i - \overline{x})(x_i - \overline{x})^T$, where $\overline{x}$ is the sample mean, and n is the number of samples of the space time process $\mathbf{X}$. The estimator of the parameters $\mathbf{T}_i$, $\mathbf{S}_i$ and $\mathbf{U}$ in (2) is performed by minimizing the following objective function

$$\|\mathbf{R} - \hat{\mathbf{R}}\|_F^2 + \beta\|\hat{\mathbf{R}} - \mathcal{R}(\mathbf{I} \otimes \mathbf{U})\|_*, \tag{3}$$

where $\mathbf{R} = \mathcal{R}(\mathbf{\Sigma}_{SCM})$ and $\mathcal{R}$ denotes the permutation rearrangement operator defined in [1, 8] which maps $pT \times pT$ matrices to $T^2 \times p^2$ matrices. The objective function is minimized over all $T^2 \times p^2$ matrices $\hat{\mathbf{R}}$ that satisfy the constraint that $\mathcal{R}^{-1}(\hat{\mathbf{R}})$ is of the form (2).

This is equivalent [2] to $\hat{\mathbf{\Sigma}} = \mathbf{I} \otimes \mathbf{U} + \mathcal{R}^{-1}(\hat{\mathbf{R}})$ where $\mathcal{R}^{-1}$ is the depermutation operator (inverse of $\mathcal{R}$) and $\hat{\mathbf{R}}$ is found by solving

$$\min_{\hat{\mathbf{R}}} ||\mathbf{M} \circ (\mathbf{R} - \hat{\mathbf{R}})||_F^2 + \beta\|\hat{\mathbf{R}}\|_*, \tag{4}$$
$$\hat{\mathbf{R}} = \sum_{i=1}^{r} t_i s_i^T,\ s.t.\ \mathbf{T}_i \text{ Toeplitz},\ \forall i,$$

where the t_i, s_i are $\mathbf{T}_i, \mathbf{S}_i$ permuted as in [8, 2], $\mathbf{M}$ is a matrix masking out the elements corresponding to the covariance diagonal, and $\circ$ denotes the elementwise product). Following the results of [7], we note that the Toeplitz constraint on $\mathbf{T}_i$ is equivalent to

$$[t_i]_k = a_{j+T}^{(i)},\ \forall k \in \mathcal{K}(j), j \in [-T+1,\ T-1], \tag{5}$$

for some vector $a^{(i)}$ where

$$\mathcal{K}(j) = \{k :\ (k-1)T + k + j \in [-T+1,\ T-1]\}. \tag{6}$$

Let

$$\tilde{t}_{j+T}^{(i)} = a_{j+T}^{(i)}\sqrt{T - |j|}, \quad j \in [-T+1,\ T-1]. \tag{7}$$

Following the method of [7], the optimization can be shown to be equivalent to

$$\min_{\tilde{\mathbf{R}}} ||\tilde{\mathbf{M}} \circ (\mathbf{B} - \tilde{\mathbf{R}})||_F^2 + \beta\|\tilde{\mathbf{R}}\|_*,\ \tilde{\mathbf{R}} = \sum_{i=1}^{r} \tilde{t}_i s_i^T, \tag{8}$$

where $\tilde{\mathbf{R}} = \mathcal{P}(\hat{\mathbf{R}})$, $\tilde{\mathbf{M}} = \mathrm{sign}(\mathcal{P}(\mathbf{M}))$, and $\mathbf{B} = \mathcal{P}(\mathbf{R})$. The operator $\mathcal{P}$ (from $T^2 \times p^2$ to $(2T-1) \times p^2$ matrices) is defined as $\tilde{\mathbf{A}} = \mathcal{P}(\mathbf{A})$ such that

$$\tilde{A}_{j+T} = \frac{1}{\sqrt{T - |j|}} \sum_{k \in \mathcal{K}(j)} A_k, \quad \forall j \in [-T+1, T-1], \tag{9}$$

where A_j is the jth row of $\mathbf{A}$.

This well-studied optimization problem (nuclear norm penalized low rank matrix approximation with missing entries) is considered in [9], where it is shown to be convex. Several fast solution methods are available, including the iterative SVD-based method of [9]. If diagonal correction is not used (i.e. $\mathbf{M}$ is all ones), the resulting optimization problem can be solved directly using the SVD [1]. The block Toeplitz diagonally corrected covariance estimate is given by

$$\hat{\mathbf{\Sigma}} = \mathbf{I} \otimes \mathbf{U} + \mathcal{R}^{-1}\left(\mathcal{P}^*\left(\tilde{\mathbf{R}}\right)\right), \tag{10}$$

where $\tilde{\mathbf{R}}$ is the minimizer of (8) and $\mathcal{P}^*$ is defined by $\mathbf{A} = \mathcal{P}^*(\tilde{\mathbf{A}})$ where

$$A_k = \frac{1}{\sqrt{T - |j|}} \tilde{A}_{j+T},\ \forall k \in \mathcal{K}(j),\ \forall j \in [-T+1,\ T-1]. \tag{11}$$

The diagonal matrix $\mathbf{U}$ is set in the next section.

3. KRONECKER PCA SHRINKAGE ESTIMATION

3.1. Diagonal Shrinkage

Diagonal shrinkage shrinks the sample covariance towards a scaled identity matrix. This improves the conditioning of the estimate, which makes the inverse of the estimate more stable.

$$\hat{\mathbf{\Sigma}} = (1 - \hat{\rho})\hat{\mathbf{\Sigma}}_{kron} + \hat{\rho}\mathbf{F} \tag{12}$$

where $\mathbf{F} = \frac{\mathrm{trace}(\hat{\mathbf{\Sigma}}_{kron})}{pT}\mathbf{I}$ and $\hat{\mathbf{\Sigma}}_{kron}$ is the DC-KronPCA estimate of the covariance.

It remains to determine the amount of shrinkage, i.e. $\hat{\rho}$. We use the Ledoit-Wolf (LW) [5] solution that asymptotically (large sample size n) minimizes Frobenius estimation

error when $\hat{\mathbf{\Sigma}}$ is the SCM. Due to the lower variance of the DC-KronPCA covariance estimate relative to the SCM, the amount of shrinkage is expected to be overestimated.

Since the LW shrinkage solution minimizes the estimator MSE it depends on the true (but unavailable) covariance, as in [5, 3] we use a plug-in estimator of $\hat{\rho}$ that uses the DC-KronPCA estimator denoted $\hat{\mathbf{\Sigma}}_{kron}$, given by (10). We call this method DC-KronPCA-LW.

3.2. Robust Shrinkage

LW methods are not designed to be robust to heavy tailed distributions [4]. We propose to modify the iterative algorithm of [4], which in turn is a modification of the Tyler iterations for ML estimation of the covariances of elliptical distributions [4] incorporating shrinkage at the end of each iteration.

Define the operator $\mathrm{KronPCA}_{\mathbf{T}}\{\cdot\}$ that returns the Toeplitz (temporal) Kronecker factor resulting from running block Toeplitz KronPCA with $r = 1$. We apply the Chen-Tyler iterations [4] modified to incorporate a projection onto covariances of the form $\mathbf{T} \otimes \mathbf{S}$. In order to guarantee positive semidefiniteness of the projection, we optimize the ML KronPCA objective function with the Toeplitz Kronecker factor held constant [8]. This problem is convex and has a closed form solution given by [8]. Our robust DC-KronPCA estimation algorithm is detailed in Algorithm 1.

Following the same rationale as in the previous subsection, we use the solution of [4] for $\hat{\rho}$, which aims to minimize the expected MSE for the SCM based method. For the plug-in covariance estimate needed in the solution [4], we use the DC-KronPCA-LW estimate.

4. SIMULATION RESULTS

We conduct numerical simulations to evaluate the relative MSE performance of the various covariance estimation algorithms considered. We first consider learning a $p = 100$, $t = 10$ multiframe covariance defined as the Kronecker product of the covariances of two AR(1) processes (equivalent to the multiplication of the processes), one across time (with AR coefficient .5) and the other space (AR coefficient .95). For various values of n, we generate n iid samples, learn the covariance, and compute the MSE. The methods compared are the SCM, the Ledoit Wolf estimator (SCM-LW), KronPCA, DC-KronPCA-LW, the robust shrinkage estimator of [4] (Chen Tyler), and our robust KronPCA shrinkage (Tyler-KronPCA). For this experiment, the Kronecker methods use $r = 1$, i.e. a single Kronecker summand in (2). The results for Gaussian measurements $\mathbf{X}$ are shown in Figure 1 (left), with all performance curves averaged over 100 trials. Figure 1 (right) shows results (averaged over 3000 trials) for the heavy tailed elliptical multivariate-t distribution given by multiplication of the AR multivariate Gaussian by $\sqrt{d/\chi^2_d}$, where $d = 3$ and χ^2_d is chi squared distributed.

Algorithm 1 Robust DC-KronPCA Shrinkage Estimation

$s_i = \frac{x_i}{\|x_i\|_2}, \forall i.$
Run Chen iterations until convergence, obtain $\hat{\mathbf{\Sigma}}$.
$\tilde{\mathbf{\Sigma}} \leftarrow \hat{\mathbf{\Sigma}}$.
while not converged **do**
  $\hat{\mathbf{T}} = \mathrm{KronPCA}_{\mathbf{T}}\{\tilde{\mathbf{\Sigma}}\}$.
  while not converged **do**
    Run Tyler iteration: $\tilde{\mathbf{\Sigma}} \leftarrow \frac{pT}{n} \sum_{i=1}^{n} \frac{s_i s_i^T}{s_i^T \hat{\mathbf{\Sigma}}^{-1} s_i}$
    $\hat{\mathbf{S}} = \frac{1}{T} \sum_{i,j=1}^{T} [\hat{\mathbf{T}}^{-1}]_{ij} \tilde{\mathbf{\Sigma}}_{ji}$.
    $\hat{\mathbf{\Sigma}} \leftarrow \hat{\mathbf{T}} \otimes \hat{\mathbf{S}}$
    $\hat{\mathbf{\Sigma}} \leftarrow (1-\rho)\frac{pT}{\mathrm{trace}(\hat{\mathbf{\Sigma}})}\hat{\mathbf{\Sigma}} + \rho \mathbf{I}$
  end while
end while
Return $\hat{\mathbf{\Sigma}}$.

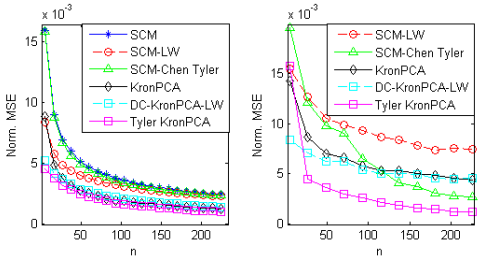

Fig. 1. Estimation of Kronecker product covariance from i.i.d. samples. Left: Normalized MSE as a function of n for Gaussian samples. Right: Normalized MSE for heavy tailed multivariate-t samples. Note the universal superiority of the Kronecker methods, the low sample improvement resulting from DC-KronPCA-LW, and the superiority of the robust shrinkage KronPCA in the heavy tailed case.

5. APPLICATION TO ANOMALY DETECTION IN WIRELESS SENSOR NETWORKS

In this section we apply our proposed spatio-temporal covariance estimators to activity detection in a wireless sensor network consisting of 14 transmit/receive Mica2 sensor nodes placed randomly in a lab at the University of Michigan [4]. All sensors transmitted asynchronously every 0.5 seconds over a period of 30 minutes, and the received signal strength (RSS) of every transmission was recorded by the remaining sensors, giving 182 measurements. Disturbances in the RSS measurements occurred when people passed through the lab at various times during this period. Ground truth was recorded using a camera [4]. The RSS measurements and ground truth are shown in [4]. As in [4], the data was detrended, the covariance learned from a subset of the data, and anomaly detection performed by thresholding the Mahanolobis distance.

For detecting anomalies in time series data, we divide the

sequence up into overlapping chunks of length T using a sliding window and determine whether each one is anomalous independently. Chunks are included in testing only if they are $> 75\%$ anomalous or nonanomalous.

The different covariance estimators described in Section 4 are used for anomaly detection. The results for the case of learning with $T = 10$ on a fixed set of 200 consecutive nonanomalous frames are shown in Figure 2. The multi-frame advantage is clearly evident (compare line labeled Chen-Tyler, $T = 1$ to the other curves).

In Figure 3, we show the magnitudes of the 19 block Toeplitz Kronecker components of the $T = 10$ SCM using the whole dataset ($n = 3000$). That KronPCA achieves higher concentration (fewer terms) than ordinary PCA is evident. We also show the learned Kronecker factors arising from robust DC-KronPCA shrinkage (for Figure 2), inverted to show partial correlation structure.

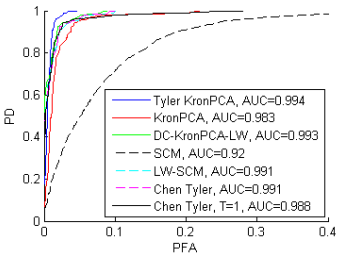

Fig. 2. RSS network anomaly ROC curves for various covariance estimation methods with $T = 10$, with the corresponding areas under the curve (AUC). Note the superiority of the Kronecker-based shrinkage methods.

6. CONCLUSION

In this work we considered the estimation of spatiotemporal covariances and their inverses in the low sample regime. In particular, we proposed a combination of Ledoit-Wolf type shrinkage techniques, Kronecker PCA, and Toeplitz Kronecker factors for temporally stationary processes. In addition, we presented a modified version of DC-KronPCA exploiting the block Toeplitz structure implied by stationarity. We validated the performance advantages of the proposed covariance estimation strategies by running several simulations and experiments.

7. ACKNOWLEDGEMENTS

This research was partially supported by ARO under grant W911NF-11-1-0391 and by AFRL under grant FA8650-07-D-1220-0006.

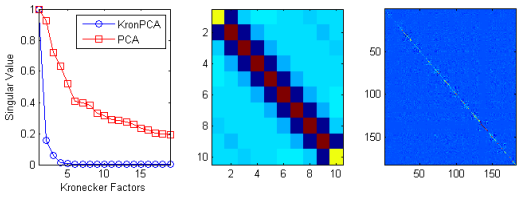

Fig. 3. Left: Normalized RMS values of (block Toeplitz) Kronecker components and first few PCA components for the SCM learned from the RSS dataset. Note that fewer Kronecker components than PCA components are required to achieve a given representation accuracy. Also shown are $\hat{\mathbf{T}}^{-1}$ (middle) and $\hat{\mathbf{S}}^{-1}$ (right), from robust DC-KronPCA shrinkage, $n = 200$.

8. REFERENCES

[1] T. Tsiligkaridis and A.O. Hero, "Covariance estimation in high dimensions via kronecker product expansions," *IEEE Trans. on Sig. Proc.*, vol. 61, no. 21, pp. 5347–5360, 2013.

[2] Kristjan Greenewald, Theodoros Tsiligkaridis, and A Hero, "Kronecker sum decompositions of space-time data," in *IEEE CAMSAP*, 2013.

[3] Y. Chen, A. Wiesel, Y. Eldar, and A. O. Hero, "Shrinkage algorithms for mmse covariance estimation," *IEEE Trans. on Sig. Proc.*, vol. 58, no. 10, pp. 5016–5029, 2010.

[4] Y. Chen, A. Wiesel, and A. Hero, "Robust shrinkage estimation of high-dimensional covariance matrices," *IEEE Trans. on Sig. Proc.*, vol. 59, no. 9, pp. 4097–4107, 2011.

[5] O. Ledoit and M. Wolf, "A well-conditioned estimator for large-dimensional covariance matrices," *Journal of multivariate analysis*, vol. 88, no. 2, pp. 365–411, 2004.

[6] A. Wiesel, O. Bibi, and A. Globerson, "Time varying autoregressive moving average models for covariance estimation," *IEEE Trans. on Sig. Proc.*, vol. 61, no. 11, pp. 2791–2801, 2013.

[7] Julie Kamm and James G Nagy, "Optimal kronecker product approximation of block toeplitz matrices," *SIAM Journal on Matrix Analysis and Applications*, vol. 22, no. 1, pp. 155–172, 2000.

[8] K. Werner, M. Jansson, and P. Stoica, "On estimation of covariance matrices with kronecker product structure," *IEEE Trans. on Sig. Proc.*, vol. 56, no. 2, pp. 478–491, 2008.

[9] Rahul Mazumder, Trevor Hastie, and Robert Tibshirani, "Spectral regularization algorithms for learning large incomplete matrices," *The Journal of Machine Learning Research*, vol. 11, pp. 2287–2322, 2010.